

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representation of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

## **IMAGES ARE BEST AVAILABLE COPY**

**As rescanning documents *will not* correct images, please do not report the images to the Image Problem Mailbox.**

**THIS PAGE BLANK (USPTO)**



⑪ Publication number : **0 570 220 A2**

⑫

## EUROPEAN PATENT APPLICATION

⑳ Application number : **93303702.0**

⑤① Int. Cl.<sup>5</sup> : **H04L 1/18**

㉔ Date of filing : **13.05.93**

③① Priority : **15.05.92 US 883639**

④③ Date of publication of application :  
**18.11.93 Bulletin 93/46**

⑧④ Designated Contracting States :  
**DE FR GB**

⑦① Applicant : **International Business Machines Corporation**  
**Old Orchard Road**  
**Armonk, N.Y. 10504 (US)**

⑦② Inventor : **Rouse, Gerald Lavelle**  
**3803 Curtis Drive**  
**Round Rock, Texas 78681 (US)**

⑦④ Representative : **Burt, Roger James, Dr.**  
**IBM United Kingdom Limited Intellectual Property Department**  
**Hursley Park**  
**Winchester Hampshire SO21 2JN (GB)**

⑤④ **Acknowledgment protocol for serial data network with out-of-order delivery.**

⑤⑦ A system and method are provided for controlling the transmission of frames or packets of data in a serial network which allows out-of-order delivery. The data frames transmitted by an initiator node to a recipient node include frame serial number or sequence count information. Acknowledgment frames, transmitted by the recipient node to indicate delivery of the data frame or packet, include matching serial number or sequence count information. According to a preferred embodiment of the present invention, the initiator node and recipient node include resources which relate the serial number or sequence count information, resources which include control information in the acknowledgment frame, and resources which determine the transmit status of all acknowledgment frames prior to that specified by serial number or sequence count in the acknowledgment frame. The status information is used at the initiator node to update the status of any data frames whose acknowledgment has not been received as a consequence of an out-of-order delivery to the recipient, a delay in processing by the recipient, a delay of the acknowledgment frame in the transmission back to the initiator node, or a complete loss of the acknowledgment frame.

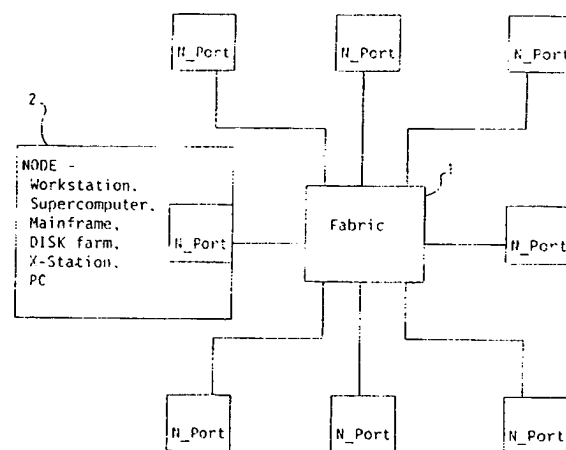


FIG. 1

The present invention relates generally to data processing systems and methods for transmitting frames or packets of information through a communication network. More particularly, the invention is directed to data processing systems and methods of operation for transmitting frames or packets of information in serial networks which utilize acknowledgments and allow out-of-order delivery of the frames or packets.

5 Numerous serial data communication protocols exist for transmitting frames or packets of information from an initiator (source) node to a recipient (destination) node in local and wide area networks. To ensure the consistency of the data at the receiving end, conventional practice involves the use of an acknowledgment transmission from the recipient node back to the initiator node upon acceptance of the information in the frame by the recipient. Though acknowledgment signals are not used as a routine in voice communication environments, 10 comprehensive acknowledgment usage is mandatory for integrity in digital data transmissions.

A commercially available example of a serial network is the fibre optical serial channel used with the RISC System/6000 workstation product (RISC System/6000 is a Trade Mark of International Business Machines Corporation) sold by IBM Corporation. Another example of such a system is defined by the Fibre Channel Standard (FCS) being developed by the American National Standards Institute (ANSI) X3T9.3 Task Group. The technology is documented in the Fibre Channel Physical and Signaling Interface (FC\_PH) Rev. 2.2 document available 15 from Global Engineering, 2805 McGaw Street, Irvine, California 92714.

A fundamental weakness of the technologies in present use and proposed in such standards arises from the need of an acknowledgment frame for each data frame transmitted by the initiator. Without specific network protocol constraints, out-of-order delivery of data frames is the serial packet network norm of operation. Even 20 if delivery of data frames or packets follows the order transmitted, there remains the possibility that the processing of the frames by the recipient to confirm adequacy of the information in the frames may differ between successive frames. The timing of the acknowledgment transmissions would follow accordingly. Similarly, the acknowledgment frame is also subject to the out-of-order delivery to the initiator as a consequence of a busy network. In some circumstances the acknowledgment frame is completely lost as a consequence of an error 25 in the network. Also possible are volitional deletions of acknowledgments by the network as a consequence of congestion or priority conditions. On the other hand, the rate of network communication is reduced significantly if the protocol requires that the data frames wait for an acknowledgment of each preceding data transmission, or that channels of the network be monopolized to ensure consistency of the data transmission and acknowledgment information.

30 In addition to acknowledgments that are lost, and as a consequence initiate error procedures within the network, standards of the type noted earlier define networks in which one or more acknowledgment frames have the misfortune of being transmitted and retransmitted into a then busy network. In each case the acknowledgment frames are returned to recipient node for later retransmission while the initiator node waits for the acknowledgment specific to the data frame.

35 Given the difficulty of handling out-of-order or lost acknowledgments, the prevailing design practice has been to require an ordered transmission in which an acknowledgment is expected for each data frame transmitted. Though the standards contemplate the use of time windows for delayed acknowledgment, the approach has proven to be relatively slow and particularly susceptible to further degradation as network congestion increases.

40 Viewed from a first aspect the present invention provides a data processing system for transmitting data in frames between nodes connected via a network, the system comprising an initiator node for transmitting data frames incorporating frame sequence information through the network, and a recipient node for receiving the data frames from the network, the recipient node including a transmitting means for transmitting an acknowledgment frame to the initiator node, the initiator node being responsive to such acknowledgment frames 45 to control the transmission of subsequent data frames, the system being characterised by: determination means in the recipient node for determining the transmission status of previous acknowledgment frames; the transmitting means being adapted to transmit an acknowledgment frame which includes the transmission status of previous acknowledgment frames to the initiator node; reception means in the initiator node for receiving the acknowledgment frame and using the transmission status in the acknowledgment frame to update the status 50 at the initiator node of data frames the subject of previous acknowledgment frames in addition to the current data frame being acknowledged.

55 Viewed from a second aspect the present invention provides a method of operating a data processing system to transmit data in frames between nodes connected via a network, the method comprising the steps of transmitting data frames incorporating frame sequence information from an initiator node through the network, and receiving at a recipient node the data frames from the network, the recipient node transmitting an acknowledgment frame to the initiator node, and the initiator node being responsive to such acknowledgment frames to control the transmission of subsequent data frames, the method being characterised by the steps of: determining, at the recipient node, the transmission status of previous acknowledgment frames; transmitting from

the recipient node an acknowledgment frame which includes the transmission status of previous acknowledgment frames to the initiator node; receiving at the initiator node the acknowledgment frame; and updating, based on the transmission status in the acknowledgment frame, the status at the initiator node of data frames the subject of previous acknowledgment frames in addition to the current data frame being acknowledged.

The present invention provides a system and method of operation which facilitates out-of-order transmission of data by the initiator, and acknowledgment information by the recipient, in frames or packets over a multi-node serial network. The architecture and protocol are implemented by system and method resources at the initiator and recipient ends of the communication path. The resources relate data frame serial number or sequence count as defined at the initiator node to corresponding serial number or sequence count in the acknowledgment information transmitted from the recipient node. The invention is characterized by its provision for controller parameter information in the acknowledgment frame, such information in preferred embodiments representing that all acknowledgments for data frames prior to the serial number or sequence count of the acknowledgment were transmitted by the recipient node. In a preferred implementation, the acknowledgment frame includes within the frame header parameter field a control bit to indicate by its state the transmission status of all preceding acknowledgment frames.

The present invention will be described further, by way of example only, with reference to an embodiment thereof as illustrated in the accompanying drawings, in which:

Figure 1 is a schematic block diagram of a network.

Figure 2 is a schematic block diagram of a fibre channel standard communication hierarchy.

Figure 3 is a schematic block diagram of a FC-2 frame format.

Figure 4 is a schematic block diagram of a frame header format.

The preferred embodiment of the present invention finds particular usage and value in systems and methods conforming to the Fibre Channel Standard (FCS) defined by the ANSI X3T9.3 Task Group. The FCS defines a serial communication architecture and protocol for a network composed of multiple nodes connected through a communication medium. Information in the form of a data frame or packet is transmitted from an initiator node to a recipient node. FCS also defines an acknowledgment protocol, in which the recipient transmits a frame or packet of information back to the initiator to indicate reception of the data frame or packet.

A high level depiction of the FCS network appears in Figure 1. Fabric 1 is the medium of communication, and broadly encompasses both optical and electronic transmission and related data manipulation resources. The fabric has connected thereto a multitude of N\_Ports 2, each comprised of initiator/recipient systems suitable to communicate over the network. Examples of node systems include workstations, super computers, disk arrays, and personal computers. Prior practices have confirmed that where switching fabric 1 allows out-of-order delivery, transmissions requiring acknowledgments are slow in the absence of a bandwidth allocation (channel monopoly) mode of operation. The poor performance for an out-of-order delivery protocol has been attributed to the fabric busy states or priority transmissions which delay mandated acknowledgment frames.

The preferred embodiment of the present invention provides out-of-order delivery of both data and acknowledgment frames while retaining network performance. The initiator node provides with each data frame transmitted to the recipient a frame serial number or sequence count for the frame. Each acknowledgment frame transmitted by the recipient, to indicate delivery of the data frame, contains the matching serial number or sequence count of the corresponding data frame which it is acknowledging. However, the acknowledgment frame also incorporates a control bit within the frame header parameter field to indicate the status at the recipient as of that data frame. The status in the acknowledgment frame indicates whether or not all acknowledgments to data frames prior to that identified by the acknowledgment frame serial number or sequence count have been transmitted by the recipient. The additional control parameter can be a single bit, indicating by binary state the composite transmission history of earlier acknowledgment frames.

A preferred implementation of the acknowledgment protocol described above draws upon the FCS network depicted in Figure 1. As defined by FCS, the channels of communication between nodes provide for any-to-any and peer-to-peer communication, as well as master-slave or initiator-target relationships. Each has resources to provide three classes of service through the fabric: Class 1 - Dedication Connection, Class 2 - Multiplex (Connectionless), and Class 3 - Datagram (Connectionless).

Each FCS channel is defined as logically composed of multiple hierarchical levels, as schematically depicted in Figure 2. The channel is not "layered" with fixed service intervals between levels, but rather, allows for a high level of chip integration of functions in successive layers. For instance, the signal protocols accomplished in layers FC-0, FC-1 and FC-2 can be implemented in a single integrated circuit chip.

FC-0 defines the physical link attributes, which include support for various data rates, single and multi-mode optical fibre, coaxial cable, and twisted pair wires. Transmitter/receiver pairs include long wave and short wave lasers, long wave LEDs, and electrical interfaces.

FC-1 uses the IBM 8B/10B transmission code to provide reliable communication. The 8B/10B code is DC-

balanced and provides special 10-bit transmission characters which facilitate byte and word alignment through the use of ordered sets. An ordered set is a four byte word which contains a K28.5 in its first position followed by three data characters. Ordered sets are used to specify primitive signals, primitive sequences and frame delimiters (Start\_of\_Frame, End\_of\_Frame).

Primitive sequences are ordered sets of digital information used to initialize or reset the fibre link. Continuous transmission of a specific primitive sequence conveys network status information. Each primitive sequence is interlocked waiting for a primitive sequence transmitted in response. Following a link initialization protocol of primitive sequences, idles (a primitive signal) are transmitted on the link. Frame transmission is accomplished by inserting a frame in the idle stream. Each frame is encapsulated between a Start\_of\_Frame (SOF) delimiter and an End\_of\_Frame (EOF) delimiter, as shown in Figure 3.

The Start\_of\_Frame delimiter delineates a frame boundary, defines a sequence boundary, and identifies the Class (1, 2, or 3) of the frame. The frame header immediately follows the SOF and contains sufficient information to control the transfer of information from a ULP in an initiator (source) N\_Port to a ULP in a recipient (destination) N\_Port. Following the frame header is the data field which may contain optional headers and payload. A 32-bit CRC checks the contents of the frame from the frame header to the end of the data field. The data field is required to be a word-multiple and may contain up to three fill bytes on the end of the data field.

The FC-2 level (Figure 2) provides the format for performing I/O operations, known as exchanges. An exchange is composed of one or more sequences flowing in one direction at any given time. Each sequence is composed of a consecutive series of data frames. As shown in Figure 3, the first data frame of a sequence is identified by a SOFix (initiate) to define the beginning of a sequence (where x specifies Class 1, 2, or 3). Other frames within the sequence are started with a SOFnx (normal).

The fields composing the frame header are shown in Figure 4. R\_CTL is used for routing control. D\_ID (Destination Identifier) and S\_ID (Source Identifier) are 24 bit fields used to identify a specific N\_Port within the address domain of the fabric. The small address was chosen to allow fast N\_Port selection within the fabric. The TYPE field specifies the upper level protocol being used. The F\_CTL (Frame Control) field provides a number of control bits which aid the signaling protocol. SEQ\_ID is a one-byte identifier to identify a specific sequence within an exchange. DF\_CTL (Data Field Control) identifies the presence of optional headers in the data field.

The SEQ\_CNT field is a frame count within the sequence. Each sequence starts with a SEQ\_CNT of 0 and is incremented by one for each frame transmitted. The OX\_ID is a 2-byte exchange identifier uniquely assigned by the initiator of the exchange. The RX\_ID is a 2-byte exchange identifier uniquely assigned by the recipient of the exchange. The parameter field is used to contain a relative offset of the payload for data frames.

In Class 1 and 2, each data frame is acknowledged by transmitting an ACK frame with a matching SEQ\_CNT. Busy and reject frames are also defined to notify the transmitter that the frame could not be delivered. A Class 1 dedicated connection is requested by transmitting a frame containing an SOFc1 delimiter. The S\_ID N\_Port waits for an ACK frame in response from the recipient N\_Port. When the ACK is received, the connection is established and the first sequence initiated. Each N\_Port may then initiate other sequences. The connection is removed by transmitting an EOFdt (disconnect-terminate) delimiter on the last ACK. A remove connection protocol is defined in order to coordinate removal of the connection.

In Class 2, each data frame is delivered as a single entity without the establishment of a dedicated connection. In Class 2, busy and reject response frames are also provided. The normal rules for originating and managing exchanges apply to both Class 1 and 2. In Class 3, each data frame is delivered on a best-effort basis without response frames.

In order to determine the operating characteristics of an N\_Port, each N\_Port performs a login sequence with the fabric and each N\_Port with which it communicates. Parameters include classes supported, number of buffers available (credit), and worldwide name. Since the number of buffers is known, each N\_Port may transmit frames up to its credit limit (M) prior to receiving an ACK frame. The ACK frame transmitted in response indicates that the buffer which received the frame is available for further frame reception. If an error is detected, the sequence is aborted and retransmitted in its entirety or from the point of failure, depending on the ULP.

In the preferred context of a fibre channel standard network operating in a Class 2 mode, the preferred embodiment of the present invention provides system and method resources at the initiator and recipient nodes to create and use control information contained in a frame header parameter field of the acknowledgment frame, such control information indicating that all prior acknowledgment frames were transmitted by the recipient. The preferred implementation utilizes a control bit in parameter field 3 of the frame header defined by the format in Figure 4. The delivery of an acknowledgment to the initiator node with the control bit set to 0 indicates that all acknowledgment frames prior to the serial number associated with that acknowledgment were in fact transmitted by the recipient node. On the other hand, the presence of a 1 in the control bit position indicates that at least one of the data frames with a serial number preceding that of the acknowledgment had,

as of the transmission time of the acknowledgment in question, either not been received, or not been processed by the recipient, been found to have an error, or not been acknowledged. Thus, the combination of the data frame serial number and control bit in the acknowledgment frame extends the status information to include both a data frame delivery status and an acknowledgment transmission history.

The use of an acknowledgment frame with a control bit indicative of prior data frame receipt and acknowledgment transmission status provides numerous network benefits. First, it permits out-of-order delivery and processing of data frames at the recipient node, since the status of the order as perceived by the initiator is reset each time the control bit conveys a historical update. Out-of-order delivery is similarly permissible for the acknowledgment frames. This is particularly important in that a network conforming to the architecture and protocol defined in the fibre channel standard is susceptible to conflicting transmissions which have priority. Such events can delay or defer for later re-transmission the affected acknowledgment frames. Furthermore, even the complete loss of an acknowledgment frame is no longer of consequence when utilizing the system and method defined by the preferred embodiment of the present invention.

In the preferred implementation, a workstation such as the RISC System/6000 with an AIX operating system, both available from IBM Corporation, executes application code at the initiator and recipient nodes to accomplish the methods defined by the invention.

The initiator is the transmitter of data frames and the receiver of acknowledgment frames. The recipient is the receiver of the data frames and transmitter of the acknowledgment frames. Both the initiator and recipient track the progress of the sequence (collection of data frames flowing in one direction) in a Sequence Status Block. In this embodiment, Fabric Busy (F\_BSY) frames are processed. In addition, Max\_Credit is used as the upper bound on the number of unacknowledged data frames that may be transmitted before receiving an acknowledgment. Credit\_Cnt is a variable in the transmitter used to count data frames which have not been acknowledged. Credit\_Cnt is incremented on data frame transmission and decremented on reception of an acknowledgment. Both F\_BSY and Credit\_Cnt are optional features. When Parm=0 in the ACK frame, all acknowledgment frames up to that FR\_CNT have been transmitted. When Parm=1 in the ACK frame, one or more previous data frames has not been acknowledged.

Pseudocode, from which source code can be derived, sets forth the preferred process to be accomplished at the initiator node. In general, the pseudocode prescribes the tracking of data frame transmissions and the receipt of related acknowledgment frames. The pseudocode begins with an initialization sequence, defines a data frame transmission tracking sequence, defines an acknowledgment receipt processing sequence, and concludes with a sequence responsive to the acknowledgment information. The variables used in the pseudocode are defined individually as follows:

FR\_CNT - the serial number transmitted in the Data frame by the Initiator (i.e. part of the frame header).

I\_FR\_CNT - the highest Data frame serial number transmitted by the Initiator which is retained in the Sequence Status Block.

ACK\_FR\_CNT - the highest serial number of any ACK frame received which is retained in the Sequence Status Block.

ACK\_R\_R - ACK\_Receive\_Register (in this example, 8 bits) indexed 0 to 7 where index value 7 (a bit) corresponds to ACK\_FR\_CNT. Index value 6 corresponds to ACK\_FR\_CNT-1 and index value 0 corresponds to ACK\_FR\_CNT-7.

CREDIT\_CNT - a value representing the number of buffers available in the Recipient for receiving Data frames.

MAX\_CREDIT - a value representing the maximum number of buffers available at the Recipient for receiving Data frames.

MISS\_ACK\_CNT - a count of ACK frames received which are out of the ACK window of processing. Credit is reclaimed after the Sequence is completed.

INITIALIZE:

```

I_FR_CNT = 0;
ACK_FR_CNT = 0;
ACK_R_R = B'11111111';
MISS_ACK_CNT = 0; /* RECLAIM AT END OF SEQUENCE */
END INITIALIZE;
```

DATA\_FRAME\_TRANSMIT:

```

IF (CREDIT_CNT < MAX_CREDIT) AND (DATA TO TRANSMIT)
THEN
  BEGIN TRANSMIT;
  BUILD FRAME HEADER;
```

```

APPEND PAYLOAD INFORMATION;
I_FR_CNT = I_FR_CNT+1;
FR_CNT = I_FR_CNT;
5   TRANSMIT FRAME;
    CREDIT_CNT = CREDIT_CNT + 1;
    END TRANSMIT;
ELSE
    WAIT FOR CREDIT;
10   WAIT FOR DATA TO TRANSMIT;
END DATA_FRAME_TRANSMIT;
F_BSY RECEIVED TO DATA FRAME:
    RETRANSMIT DATA FRAME;
END F_BSY TO DATA FRAME;
15  ACK_RECEPTION:
    IF (FR_CNT = ACK_FR_CNT)
        THEN DETECT DUPLICATE FRAME ERROR;
    IF (FR_CNT < ACK_FR_CNT)
        THEN BEGIN;
20      IF (ACK_R_R = B'11111111')
          THEN DISCARD ACK;
        ELSE BEGIN;
          BITNUM = ACK_FR_CNT - FR_CNT;
          SETBIT = 7 - BITNUM;
25      SET ACK_R_R (BITNUM) = 1;
          CREDIT_CNT = CREDIT_CNT - 1;
          IF (PARM = 0)
            THEN DO I = BITNUM TO 0 BY -1;
              IF ACK_R_R (I) = 0;
30              THEN BEGIN;
                CREDIT_CNT = CREDIT_CNT - 1;
                ACK_R_R (I) = 1;
                END BEGIN;
            END DO;
          ELSE;
35          END BEGIN;
        END BEGIN;
    ELSE;
    IF (FR_CNT > ACK_FR_CNT)
40    THEN BEGIN;
      GOBIT = 1;
      DO WHILE (GOBIT=1);
      IF (ACK_R_R (0) = 0)
        THEN SET MISS_ACK_CNT = MISS_ACK_CNT + 1;
45      ELSE;
        BEGIN;
          /* SHIFT ACK_R_R LEFT ONE BIT */
          ACK_R_R (0) = ACK_R_R (1);
          ACK_R_R (1) = ACK_R_R (2);
50      ACK_R_R (2) = ACK_R_R (3);
          ACK_R_R (3) = ACK_R_R (4);
          ACK_R_R (4) = ACK_R_R (5);
          ACK_R_R (5) = ACK_R_R (6);
          ACK_R_R (6) = ACK_R_R (7);
55      IF (PARM = 0)
        THEN BEGIN;
          ACK_R_R (7) = 1;
          CREDIT_CNT = CREDIT_CNT - 1;
          END BEGIN;

```



```

ELSE ACK_R_R (7) = 0; /* PARM = 1 */
ACK_FR_CNT = ACK_FR_CNT + 1;
IF (FR_CNT = ACK_FR_CNT) AND (ACK_R_R (7) = 1)
5 THEN BEGIN;
  GOBIT = 0;
  EXIT;
  END BEGIN;
ELSE;
10 IF (FR_CNT = ACK_FR_CNT) AND (ACK_R_R (7) = 0)
  THEN BEGIN;
    CREDIT_CNT = CREDIT_CNT - 1;
    ACK_R_R (7) = 1;
    GOBIT = 0;
15 EXIT;
    END BEGIN;
  ELSE;
    END BEGIN;
  END DO_WHILE;
20 END BEGIN;
  ELSE;

```

END ACK\_RECEPTION;

The system and process activities at the recipient node generally track the data frame reception and the acknowledgment frame transmission operations. The pseudocode begins with an initialization sequence, defines operations related to the reception of a data frame, defines processes associated with creating a frame header for the acknowledgment frame, undertakes the transmission of the acknowledgment frame, and defines a response to a busy state of the network fibre. The variables used in the pseudocode set forth below for the recipient node are defined as follows:

FR\_CNT - the serial number received in the Data frame by the Recipient (i.e. part of the frame header) and transmitted in the ACK frame acknowledging the Data frame specified by FR\_CNT.

DF\_FR\_CNT - the highest Data frame serial number received by the Recipient which is retained in the Sequence Status Block.

DF\_R\_R - Data frame \_Receive\_Register (in this example, 8 bits) indexed 0 to 7 where index value 7 (a bit) corresponds to DF\_FR\_CNT. Index value 6 corresponds to DF\_FR\_CNT-1 and index value 0 corresponds to DF\_FR\_CNT-7. This register is used to detect missing Data frames.

ACK\_FR\_CNT - the highest serial number of any ACK frame transmitted which is retained in the Sequence Status Block.

ACK\_T\_R - ACK\_Transmit\_Register (in this example, 8 bits) indexed 0 to 7 where index value 7 (a bit) corresponds to ACK\_FR\_CNT. Index value 6 corresponds to ACK\_FR\_CNT-1 and index value 0 corresponds to ACK\_FR\_CNT-7.

CREDIT\_CNT - a value representing the number of buffers available in the Recipient for receiving Data frames.

BITNUM - a register used in calculations to set bits in a register.

SETBIT - a register used in calculations to set bits in a register.

GOBIT - a single bit variable used to control a DO\_WHILE loop.

```

INITIALIZE:
  DF_R_R = B'11111111';
  ACK_FR_CNT = 0;
  ACK_T_R = B'11111111';
50 END INITIALIZE;
DATA_FRAME_RECEPTION:
  IF (BUFFER AVAILABLE)
    THEN
      BEGIN STORE_BUFFER;
55 IF (FIRST FRAME OF SEQUENCE)
        THEN BEGIN;
          DR_FR_CNT = FR_CNT;
          EXIT;
          END BEGIN;

```

```

ELSE;
IF (FR_CNT < DF_FR_CNT)
THEN BEGIN;
5   BITNUM = DF_FR_CNT - FR_CNT;
   SETBIT = 7 - BITNUM;
   SET DF_R_R (BITNUM) = 1;
   END BEGIN;
ELSE BEGIN; /* DF_FR_CNT < FR_CNT */
10  GOBIT = 1;
   DO WHILE (GOBIT=1);
     IF (DF_R_R(0)=0)
       THEN SET MISSING_FRAME_ERROR;
       ELSE BEGIN;
15     /* SHIFT DF_R_R LEFT ONE BIT */
       DF_R_R (0) = DF_R_R (1);
       DF_R_R (1) = DF_R_R (2);
       DF_R_R (2) = DF_R_R (3);
       DF_R_R (3) = DF_R_R (4);
20     DF_R_R (4) = DF_R_R (5);
       DF_R_R (5) = DF_R_R (6);
       DF_R_R (6) = DF_R_R (7);
       DF_R_R (7) = 0;
       DF_FR_CNT = DF_FR_CNT + 1;
25     IF (FR_CNT = DF_FR_CNT)
       THEN BEGIN;
         SET DF_R_R (7) = 1;
         GOBIT = 0;
         EXIT;
30     END BEGIN;
       ELSE GOBIT = 1;
       END BEGIN;
     END DO_WHILE;
     PASS TO BUFFER MANAGEMENT;
35     BUFFER MANAGEMENT INDICATES BUFFER EMPTIED
     END STORE_BUFFER;
ELSE
  TRANSMIT P_BSY;
  END DATA_FRAME_RECEPTION;
40  ACK_TRANSMISSION (BUFFER EMPTIED);
  USE DATA FRAME HEADER TO CONSTRUCT ACK FRAME;
  IF (FR_CNT > ACK_FR_CNT)
  THEN BEGIN;
    GOBIT = 1;
45    DO WHILE (GOBIT=1);
      IF (ACK_T_R(0)=0)
      THEN SET PARM = 1;
      ELSE BEGIN;
        /* SHIFT ACK_T_R LEFT ONE BIT */
50        ACK_T_R (0) = ACK_T_R (1);
        ACK_T_R (1) = ACK_T_R (2);
        ACK_T_R (2) = ACK_T_R (3);
        ACK_T_R (3) = ACK_T_R (4);
        ACK_T_R (4) = ACK_T_R (5);
55        ACK_T_R (5) = ACK_T_R (6);
        ACK_T_R (6) = ACK_T_R (7);
        ACK_T_R (7) = 0;
        ACK_FR_CNT = ACK_FR_CNT + 1;
        IF (FR_CNT = ACK_FR_CNT)

```

```

        THEN BEGIN;
          SET ACK_T_R (7) = 1;
          GOBIT = 0;
5          EXIT;
        END BEGIN;
      ELSE GOBIT = 1;
    END BEGIN;
  END DO_WHILE;
10  DO I = 7 TO 0 BY -1;
    IF (ACK_T_R (I) = 0)
      THEN BEGIN;
        PARM = 1;
        EXIT;
15      END BEGIN;
    ELSE PARM = 0;
  END DO;
  TRANSMIT ACK FRAME;
  END BEGIN;
20 ELSE;
  IF (FR_CNT < ACK_FR_CNT)
    THEN BEGIN;
      BITNUM = ACK_FR_CNT - FR_CNT;
      SETBIT = 7 - BITNUM;
25      SET ACK_T_R (BITNUM) = 1;
      DO I = BITNUM TO 0 BY -1;
        IF ACK_T_R (I) = 0;
          THEN BEGIN;
            PARM = 1;
30            EXIT;
          END BEGIN;
        ELSE PARM = 0;
      END DO;
      TRANSMIT ACK FRAME;
35      EXIT;
    END BEGIN;
  END ACK_TRANSMISSION;
  F_BSY_TO_ACK:
  BITNUM = 0;
40  DO I = 7 TO 0 BY -1; /* COUNT ACKS NOT TRANSMITTED */
    IF (ACK_T_R (I) = 0)
      THEN BITNUM = BITNUM + 1;
    ELSE;
  END DO;
45  IF (BITNUM = MAX_CREDIT) /* INITIATOR OUT OF CREDIT */
    THEN RETRANSMIT ACK /* THEN RETRANSMIT ACK */
    ELSE DISCARD F_BSY TO ACK;
  END F_BSY_TO_ACK;

```

50

### Claims

1. A data processing system for transmitting data in frames between nodes connected via a network, the system comprising an initiator node for transmitting data frames incorporating frame sequence information through the network, and a recipient node for receiving the data frames from the network, the recipient node including a transmitting means for transmitting an acknowledgment frame to the initiator node, the initiator node being responsive to such acknowledgment frames to control the transmission of subsequent data frames, the system being characterised by:
  - determination means in the recipient node for determining the transmission status of previous ac-

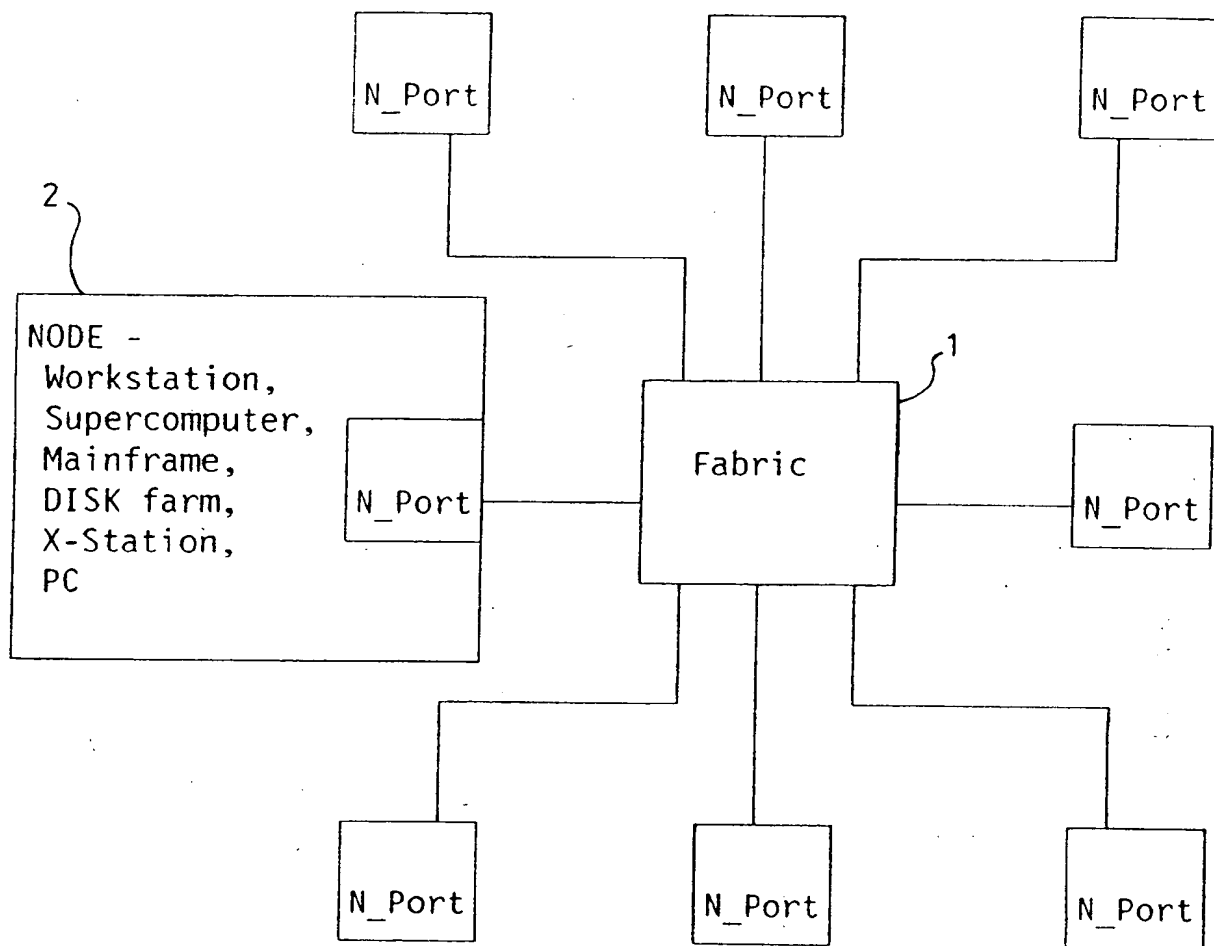
55

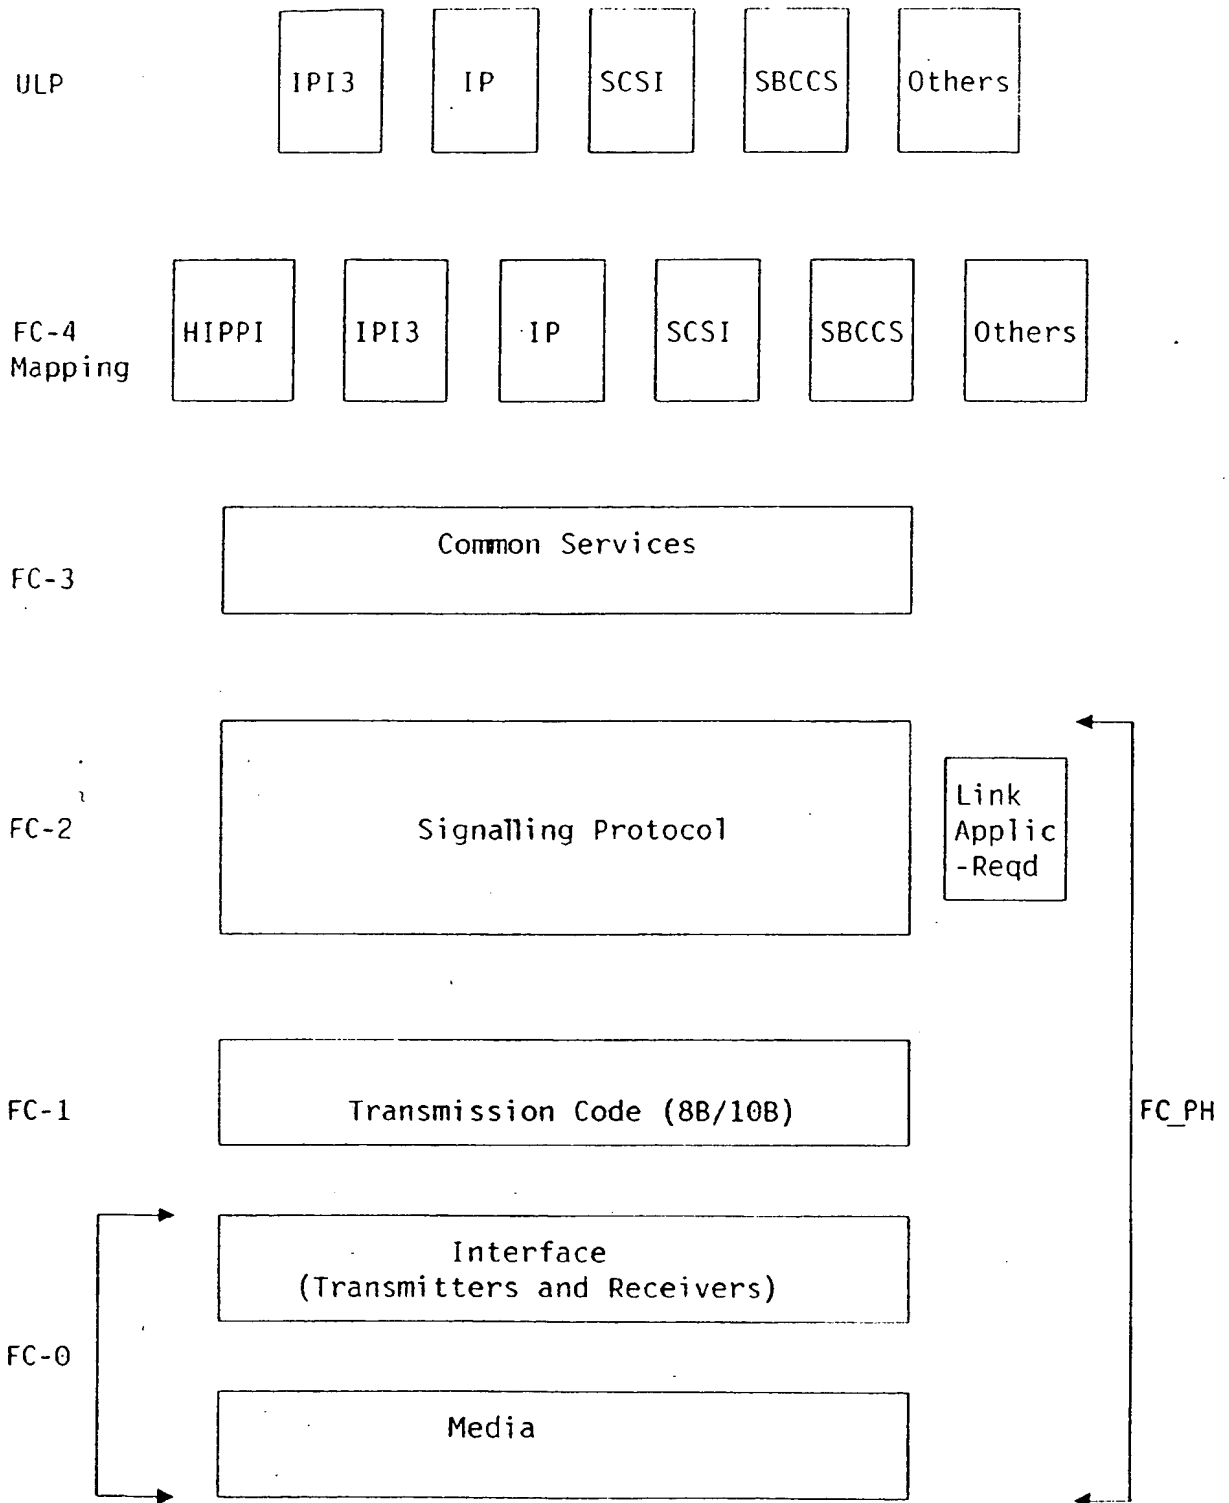
knowledgegment frames;

the transmitting means being adapted to transmit an acknowledgment frame which includes the transmission status of previous acknowledgment frames to the initiator node;

5 reception means in the initiator node for receiving the acknowledgment frame and using the transmission status in the acknowledgment frame to update the status at the initiator node of data frames the subject of previous acknowledgment frames in addition to the current data frame being acknowledged.

- 10 2. A data processing system as claimed in claim 1, wherein the frame sequence information is a serial number or frame count attributed to the frame of data.
3. A data processing system as claimed in claim 1 or claim 2, wherein the status of previous acknowledgment frames indicates whether all preceding acknowledgment frames have been transmitted.
- 15 4. A data processing system as claimed in any of claims 1, 2 or 3, wherein the reception means cancels any error or retransmit conditions for data frames previously acknowledged.
5. A data processing system as claimed in any preceding claim wherein the frame data is transmitted over a serial data network susceptible to frame loss, delay or rearranged order upon delivery.
- 20 6. A method of operating a data processing system to transmit data in frames between nodes connected via a network, the method comprising the steps of transmitting data frames incorporating frame sequence information from an initiator node through the network, and receiving at a recipient node the data frames from the network, the recipient node transmitting an acknowledgment frame to the initiator node, and the initiator node being responsive to such acknowledgment frames to control the transmission of subsequent data frames, the method being characterised by the steps of:
  - 25 determining, at the recipient node, the transmission status of previous acknowledgment frames;
  - transmitting from the recipient node an acknowledgment frame which includes the transmission status of previous acknowledgment frames to the initiator node;
  - receiving at the initiator node the acknowledgment frame; and
  - 30 updating, based on the transmission status in the acknowledgment frame, the status at the initiator node of data frames the subject of previous acknowledgment frames in addition to the current data frame being acknowledged.
- 35 7. A method as claimed in claim 6, wherein the frame sequence information is a serial number or frame count attributed to the frame of data.

FIG. 1



**FIG. 2**


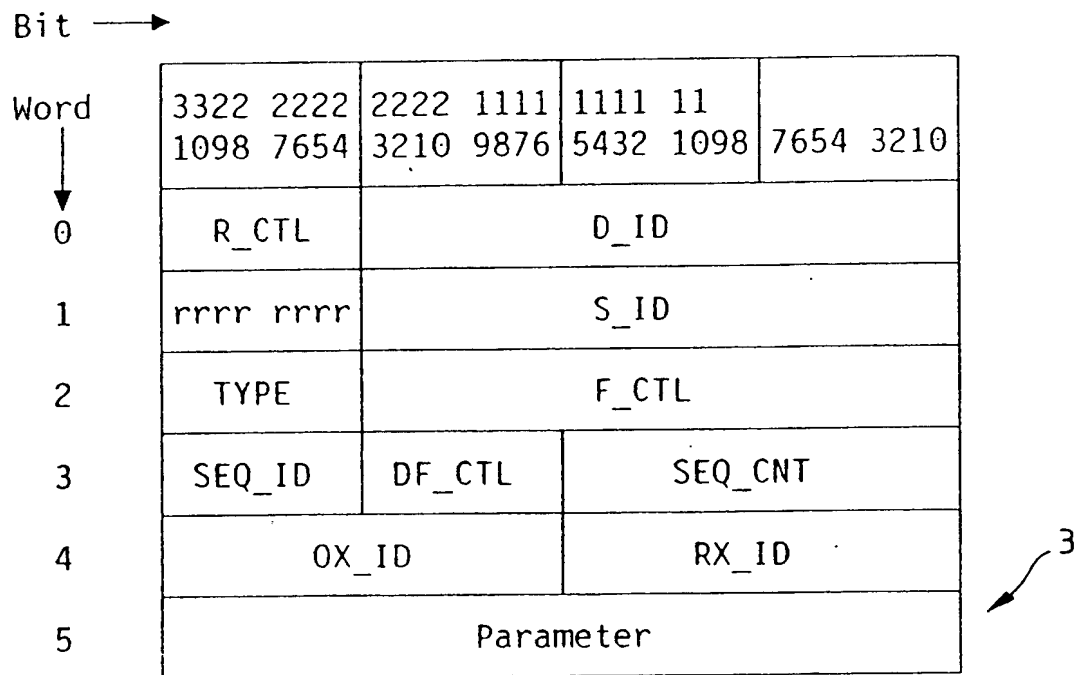


FIG. 3

FIG. 4



(19)



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11)

**EP 0 570 220 A3**

(12)

**EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:  
15.05.1996 Bulletin 1996/20

(51) Int Cl.<sup>6</sup>: **H04L 1/18**, H04L 12/56,  
H04L 29/06, H04L 1/16

(43) Date of publication A2:  
18.11.1993 Bulletin 1993/46

(21) Application number: **93303702.0**

(22) Date of filing: **13.05.1993**

(84) Designated Contracting States:  
**DE FR GB**

(72) Inventor: **Rouse, Gerald Lavelle**  
**Round Rock, Texas 78681 (US)**

(30) Priority: **15.05.1992 US 883639**

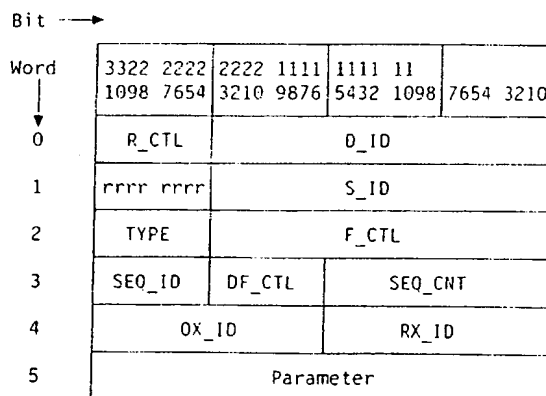
(74) Representative: **Burt, Roger James, Dr.**  
**IBM United Kingdom Limited**  
**Intellectual Property Department**  
**Hursley Park**  
**Winchester Hampshire SO21 2JN (GB)**

(71) Applicant: **International Business Machines Corporation**  
**Armonk, N.Y. 10504 (US)**

(54) **Acknowledgment protocol for serial data network with out-of-order delivery**

(57) A system and method are provided for controlling the transmission of frames or packets of data in a serial network which allows out-of-order delivery. The data frames transmitted by an initiator node to a recipient node include frame serial number or sequence count information. Acknowledgment frames, transmitted by the recipient node to indicate delivery of the data frame or packet, include matching serial number or sequence count information. According to a preferred embodiment of the present invention, the initiator node and recipient node include resources which relate the serial number or sequence count information, resources which include

control information in the acknowledgment frame, and resources which determine the transmit status of all acknowledgment frames prior to that specified by serial number or sequence count in the acknowledgment frame. The status information is used at the initiator node to update the status of any data frames whose acknowledgment has not been received as a consequence of an out-of-order delivery to the recipient, a delay in processing by the recipient, a delay of the acknowledgment frame in the transmission back to the initiator node, or a complete loss of the acknowledgment frame.

**FIG. 4****EP 0 570 220 A3**



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 93 30 3702

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.CLS)
X A	GB-A-2 229 896 (ROSEMOUNT INC) * abstract * * page 1, line 24 - page 4, line 22; figures 2-4 * * page 5, line 2 - page 7, line 2 * * claims *	1,2,5-7 3,4	H04L1/18 H04L12/56 H04L29/06 H04L1/16
A	US-A-5 084 877 (NETRAVALI ET AL.) * abstract * * column 2, line 56 - line 68 * * column 4, line 50 - column 5, line 34 * * claims; figures *	1-7	
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 21 March 1996	Examiner Gries, T
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons &amp; : member of the same patent family, corresponding document</p>			

EPO FORM 1503 (03/92) (P/ACT/CH)